arXiv:1411.5867v2 [cs.DS] 4 Apr 2015

Planar Reachability in Linear Space and Constant Time

Jacob Holm and Eva Rotenberg and Mikkel Thorup*†

University of Copenhagen (DIKU),
jaho@di.ku.dk, roden@di.ku.dk, mthorup@di.ku.dk

April 7, 2015

Abstract

We show how to represent a planar digraph in linear space so that reachability queries can be answered in constant time. The data structure can be constructed in linear time. This representation of reachability is thus optimal in both time and space, and has optimal construction time. The previous best solution used $O(n \log n)$ space for constant query time [Thorup FOCS'01].

*Research partly supported by Thorup's Advanced Grant from the Danish Council for Independent Research under the Sapere Aude research career programme.

†Research partly supported by the FNU project AlgoDisc - Discrete Mathematics, Algorithms, and Data Structures.

1 Introduction

Representing reachability of a directed graph is a fundamental challenge. We want to represent a digraph $G = (V, E)$, $n = |V|$, $m = |E|$, so that we for any vertices u and w can tell if u reaches v, that is, if there is a dipath from u to v. There are two extreme solutions: one is to just store the graph, as is, using $O(m)$ words of space and answering reachability queries from scratch, e.g., using breadth-first-search, in $O(m)$ time. The other is to store a reachability matrix using n^2 bits and then answer reachability queries in constant time. Thorup and Zwick [20] proved that there are graphs classes such that any representation of reachability needs $\Omega(m)$ bits. Also, Pătraşcu [16] has proved that there are directed graphs with $O(n)$ edges where constant time reachability queries require $n^{1+\Omega(1)}$ space. Thus, for constant time reachability queries to a general digraph, all we know is that the worst-case space is somewhere between $\Omega(m + n^{1+\Omega(1)})$ and n^2 bits.

The situation is in stark contrast to the situation for undirected/symmetric graphs where we can trivially represent reachability queries on $O(n)$ space and constant time, simply by enumerating the connected components, and storing with each vertex the number of the component it belongs to. Then u reaches v if and only if the have the same component number.

In this paper we focus on the planar case, which feels particularly relevant when you live on a sphere. For planar digraphs it is already known that we can do much better than for general digraphs. Back in 2001, Thorup [19] presented a reachability oracle for planar digraphs using $O(n \lg n)$ space for constant query time, or linear space for $O(\log n)$ query time. In this paper, we present the first improvement; namely an $O(n)$ space reachability oracle that can answer reachability queries in constant time. Note that this bound is asymptotically optimal; even to distinguish between the subclass of directed paths of length n, we need $\Omega(n \log n)$ bits. Our oracle is constructed in linear time.

Computational model The computational model for all upper bounds is the word RAM, modelling what we can program in a standard programming language such as C [12]. A word is a unit of space big enough to fit any vertex identifier, so a word has $w \geq \lg n$ bits, and word operations take constant time. Here $\lg = \log_2$. In our upper bounds, we limit ourselves to the *practical RAM* model [14], which is a restriction of the word RAM to the standard operations on words available in C that are AC0. This includes indexing arrays as needed just to store a reachability matrix with constant time access, but excludes e.g. multiplication and division. Thus, unless otherwise specified, we measure *space* as the number of words used and *time* as the number of word operations performed.

The $\Omega(m + n^{1+\Omega(1)})$ space lower bound from [16] for general graphs is in the cell-probe model subsuming the word RAM with an arbitrary instruction set.

Other related work Before [19], the best reachability oracles for general planar digraphs were distance oracles, telling not just if u reaches w, but if so, also the length of the shortest dipath from u to w [3–5]. For such planar distance oracles, the best current time-space trade-off is $\tilde{O}(n/\sqrt{s})$ time for any $s \in [n, n^2]$ [15].

The construction of [19] also yields approximate distance oracles for planar digraphs. With edge weights from $[N]$, $N \leq 2^w$, distance queries where answered within a factor $(1 + \epsilon)$ in $O(\log\log(Nn) + 1/\varepsilon)$ time using $O(n(\log n)(\log(Nn))/\varepsilon)$ space. These bounds have not been improved.

For the simpler case of undirected graphs, where reachability is trivial, [13, 19] provides a more efficient $(1 + \varepsilon)$-approximate distance queries for planar graphs in $O(1/\varepsilon)$ time and $O(n(\log n)/\varepsilon)$ space. In [10] it was shown that the space can be improved to linear if the query time is increased to $O((\log n)^2/\varepsilon^2)$. In [11] it was shown how to represent planar graphs with bounded weights using $O(n \log^2((\log n)/\varepsilon) \log^*(n) \log\log(1/\varepsilon))$ space and answering $(1 + \varepsilon)$ approximate distance queries in $O((1/\varepsilon)\log(1/\varepsilon)\log\log(1/\varepsilon)\log^*(n) + \log\log\log n))$ time. Using $\bar{O}$ to suppress factors of $O(\log\log n)$ and $O(\log(1/\varepsilon))$, these bounds reduce

to $\bar{O}(n)$ space and $\bar{O}(1/\varepsilon)$ time. This improvement is similar in spirit to our improvement for reachability in planar digraphs. However, the techniques are entirely different.

There has also been work on special classes of planar digraphs. In particular, for a planar s-t-graph, where all vertices are on dipaths between s and t, Tamassia and Tollis [17] have shown that we can represent reachability in linear space, answering reachability queries in constant time. Also, [4,6,7] presents improved bounds for planar exact distance oracles when all the vertices are on the boundary of a small set of faces.

Techniques We will develop our linear space constant query time reachability oracles by considering more and more complex classes of planar digraphs. We make reductions from $i+1$ to i in the following:

1. Acyclic planar s-t-graph; $\exists(s,t)$, such that all vertices are reachable from s and may reach t. [17]
2. Acyclic planar single-source graph; $\exists s$, such that all vertices are reachable from s. See Section 3.
3. Acyclic planar In-Out graph; $\exists s$ such that all vertices with out-degree 0 are reachable from s. See Section 4
4. Any acyclic planar graph. The reduction to acyclic planar In-Out graphs from general acyclic planar graphs is known. [19]
5. Any planar graph. The reduction to acyclic planar graphs is well-known. Using the depth first search algorithm by Tarjan [18], we can contract each strongly connected component to get an acyclic planar graph. Vertices in the same strongly connected component can always reach each other, and vertices in distinct strongly connected components can reach each other if the corresponding vertices in the contracted graph can.

The most technically involved step is the reduction from single-source graph to s-t-graph. As in [19], we use separators to form a tree over a partitioning of the vertices of the graph. However, in [19], the *alternation number*; the number of directed segments in the frame that separates a child from its parent (see Section 2), needs only be a constant number. In contrast, it is a crucial part of our construction that the alternation number, which must be even, is at most 4. Also, in our data structure, paths cannot go upward in the rooted tree, whereas there is no such restriction in [19]. These two features let us use a level ancestor -like algorithm to quickly calculate the best ≤ 4 vertices in a given tree-node that can reach a given vertex v. Each component is an s-t-graph, and v can be reached by some u in the ancestral component if and only if u can reach at least one of these best ≤ 4 vertices.

2 Preliminaries

For a vertex v at depth d in a rooted forest T and an integer $0 \leq i \leq d$, the i'th *level ancestor* of v in T is the ancestor to v in T at depth i. For two nodes x, y in a rooted tree, let $x \preceq y$ denote that x is an ancestor to y, and $x \prec y$ that x is a proper ancestor to y.

We say a graph is *plane*, if it is embedded in the plane, and denote by π_v the permutation of edges around v. Given a plane graph, (G,π), we may introduce *corners* to describe the incidence of a vertex to a face. A vertex of degree n has n corners, where if $\pi_v((v,u)) = (v,w)$, and the face f is incident to (v,u) and (v,w), then there is a corner of f incident to v between (v,u) and (v,w). We denote by $V[X]$ and $E[X]$ the vertices and edges, of some (not necessarily induced) subgraph X. Given a subgraph H of a planar embedded graph G, the faces of H define *superfaces* of those of G, and the faces of G are *subfaces* of those of H. Similarly for corners. Note that the faces of H correspond to the connected components of $G^* \setminus H$. The super-corners incident to v correspond to a set of consecutive corners in the ordering around v.

In a directed graph, we may consider the boundary of a face in some subgraph, H. A corner of a face f of H is a *target* for f if it lies between ingoing edges (u, v) and (w, v), and *source* if it lies between outgoing edges (v, u) and (v, w). We say the face boundary has *alternation number* $2a$ if it has a source and a target corners. When a face boundary has alternation number $2a$, we say it consists of $2a$ *disegments* (*di*rected segments), associated with the directed paths from source to target. We associate with each disegment the total ordering stemming from reachability of vertices on the path via the path, and by convention we set $\mathrm{succ}(t, S) = \perp$ for a target vertex t on the disegment. Given a set of edges $S \subset E$, we denote by $\mathrm{init}(S)$ the set of inital vertices, $\mathrm{init}(S) = \{u | (u, v) \in S\}$. Given a connected planar graph with a spanning tree T, the edges $T^* := E \setminus T$ form a spanning tree for the dual graph. We call the pair (T, T^*) a tree-cotree decomposition of the graph, referring to T and T^* as tree and cotree.

When u can reach v we write $u \rightsquigarrow v$. An s-t-graph is a graph with special vertices s, t such that $s \rightsquigarrow v$ and $v \rightsquigarrow t$ for all vertices v. We say a graph is a *truncated s-t-graph* if it is possible to add vertices s, t to obtain an s-t-graph, without violating the embedding. In an s-t-graph, all faces has alternation number 2.

3 Acyclic planar single-source digraph

Given a global source vertex s for the planar digraph, we wish to make a data structure for reachability queries. We do this by reduction to the s-t-case. A tree-like structure with truncated s-t-graphs as nodes is obtained by recursively choosing a face f wisely, and then letting vertices that can reach vertices on f belong to this node, and partitioning all other vertices among the descendants of this node. As we shall see in Section 3.1, this can be done in such a way that we obtain logarithmic height and such that the border between a node and its ancestors is a cycle of alternation number at most 4. We call this the *frame* of the node.

We always choose the truncated s-t-graph maximally, such that once a path crosses a frame, it does not exit the frame again. Thus, for u to reach v, u has to lie in a component which is ancestral to that of v, and since the alternation number of any frame between those two component is at most 4, the path could always be chosen to use one of the at most 4 different "best" vertices for reaching v on that frame. Thus, the idea is to do something inspired by level ancestry to find those "best" vertices in u's component. We handle the case of frames with alternation number 2 in Section 3.3. Frames with alternation number 4 are similar but more involved, and the details are found in Section 3.4.

Definition 3.1. Given a graph $G = (V, E)$, a subgraph $G' = (V', E')$ is *backward closed* if $\forall (u, v) \in E : v \in V' \implies (u, v) \in E'$.

Definition 3.2. The *backward closure* of a face f, denoted $\mathrm{bc}(f)$ is the unique smallest backward closed graph that contains all the vertices incident to f.

Definition 3.3. Let $G = (V, E)$ be an acyclic single-source plane digraph, and let $G^* = (V^*, E^*)$ be its dual. An *s-t-decomposition* of G is a rooted tree where each node x is associated with a face $f_x \in V^*$ and subgraphs $G^*_x \subseteq G^*$ and $C_x \subseteq S_x \subseteq G$ such that:

- f_x is unique ($f_x \neq f_y$ for $x \neq y$).
- S_x is $\mathrm{bc}(f_x)$ if x is the root, and $\mathrm{bc}(f_x) \cup S_y$ if x is a child of y.
- C_x is $\mathrm{bc}(f_x)$ if x is the root, and $\mathrm{bc}(f_x) \setminus S_y$ if x is a child of y.

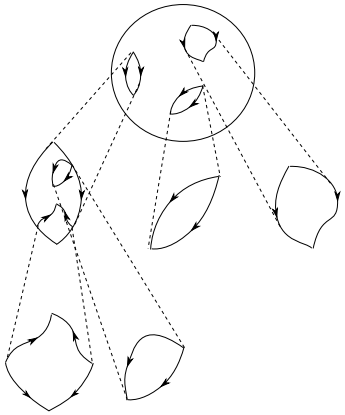

Figure 1: A tree of truncated s-t-graphs, each child contained in a face-cycle of its parent.

- G^*_x is the subgraph of G^* induced by $\{f_z |\ z \text{ is a descendent of } x\}$. Furthermore, if x is a child of y we require that G^*_x is the connected component of $G^* \setminus E^*[S_y]$ containing f_x.

If x is a child of y, x has a *parent frame* $F_x \subseteq S_y$ and a set of *down-edges* $E_x \subseteq E$ such that:

- F_x is the face cycle in S_y that corresponds to G^*_x.
- E_x is the set of edges (w, w') such that $w \in V[F_x]$ and $w' \in V[C_z]$ for some descendant z of x.

An s-t-decomposition is *good* if the tree has height $\mathcal{O}(\log n)$ and each frame has alternation number 2 or 4.

The name s-t-decomposition is chosen based on the following

Lemma 3.4. Each vertex of G is in exactly one C_x, and each C_x is a truncated s-t-graph.

Proof. If x is the root, $C_x = \mathrm{bc}(f_x)$ and this is clearly a truncted s-t-graph. Otherwise let y be the parent of x. Then $S_x = \mathrm{bc}(f_x) \cup S_y$, is backward-closed and therefore contains s. Contracting S_y in that graph to a single vertex s' gives a single-source graph S_x/S_y with s' as the source. Adding a dummy target t' in f_x results in an s-t-graph $(S_x/S_y) \cup \{t'\}$. Thus, S_x/S_y is a truncated s-t-graph, and since $C_x = \mathrm{bc}(f_x) \setminus S_y = S_x \setminus S_y = (S_x/S_y) \setminus \{s'\}$ so is C_x.

Let v be a vertex, let I be the set of all nodes in the s-t-decomposition whose associated faces $\{f_x\}_{x \in I}$ are reachable from v, and let $N = \mathrm{lca}(I)$. We now show that v lies in C_N and only in C_N. To see that $v \in C_N$, note that $v \in S_x$ for all $x \in I$, but then $v \in \bigcap_{x \in I} S_x = S_N$. But $v \notin S_a$ for any ancestor a of N by definition of lca, and thus, $v \notin S_y$ for the parent y of N, entailing $v \in S_N \setminus S_y = C_N$. We have now seen that $v \in C_N$ and that $v \notin C_x$ when $x \prec N$ or $N \prec x$. To see that $v \notin C_x$ for any unrelated $x \neq N$, note the following: if x has no descendants in I, then $v \notin V[S_x]$ since all vertices reachable from v lie on some face. Thus, $v \notin C_x \subseteq S_x$.

□

Theorem 3.5. *Any acyclic single-source plane digraph has a good s-t-decomposition.*

We defer the proof to section 3.1. The reason for studying s-t-decompositions in the context of reachability is the following

Lemma 3.6. If $u \rightsquigarrow v$ where $u \in C_x$ and $v \in C_y$ then either $x = y$ or x has a child z that is ancestor to y such that any $u \rightsquigarrow v$ path contains a vertex in F_z.

Proof. Note that whenever $w \rightsquigarrow w'$ with $w' \in C_a$, w must belong to an ancestor of a, since $w \in bc(f_a)$. Thus, x is an ancestor of y, which means that either $x = y$ or x has a child z that is an ancestor to y. But then either w lies on F_x, or F_x is a cycle separating w from w'. In either case, a path from w to w' must contain a vertex on F_x. □

Since (by theorem 3.5) we can assume the alternation number is at most 4, this reduces the reachability question to the problem of finding the at most 4 "last" vertices on $F_z \cap C_x$ that can reach v and then checking in C_x if u can reach either of them. In section 3.3 we will show how to do this efficiently when F_z is *a 2-frame*, that is, has alternation number 2, and in section 3.4 we will extend this to the case when F_z is *a 4-frame*, that is, has alternation number 4.

Theorem 3.7. *There exists a practical RAM data structure that for any planar digraph with n vertices uses $O(n)$ words of $O(\log n)$ bits and can answer reachability queries in constant time. The data structure can be built in linear time.*

Proof. First, build a good s-t-decomposition of G. Such a decomposition exists (Lemma 3.5) and can be built in linear time (Lemma 3.15). Adding DFS pre- and postorder numbers to each node in the tree lets us discover the ancestry relationship between any two vertices in constant time. Then, calculate the structures described in Section 3.3 (in particular $d_2[]$) and Section 3.4 ($c[]$ and $d[]$).

To answer $\mathrm{reachable}(u, v)$, there are the following cases. Let $x = c[u]$ and $y = c[v]$.

1. If $x \not\preceq y$, then u cannot reach v.
2. If $x = y$, then the answer is given by the s-t-graph labelling of C_x from [17].
3. If $x \prec y$ and $d_2[u] == d_2[v]$ there are no 2-frames separating u and v, but since $x \prec y$ there are 4-frames. Let $i = d[u]$, then by 3.51 we can in constant time compute $l_i^0(v)$, $r_i^0(v)$, $l_i^1(v)$, and $r_i^1(v)$. If u can reach any of them, them u can reach v, otherwise no.
4. Otherwise $x \prec y$ and $d_2[u] < d_2[v]$ and there is a 2-frame separating u and v. Let $i = d_2[u]$, then by 3.33 we can in constant time compute $l_i(v)$ and $r_i(v)$. If u can reach any of them, then u can reach v, otherwise no.

Note that the recursive calls in step 3 only leads to questions of type < 3, and similarly the recursive calls in step 4 only leads to questions of type < 4. Thus any query uses case 3 at most twice and case $1 + 2$ at most 8 times. Thus we use only constant time per query. □

A consequence of our construction which might be of independent interest is the following:

Theorem 3.8. *If a planar digraph G admits an s-t-decompostion of height h where all frames have alternation number 2 and 4, there exists an $O(h \log n)$ bit labelling scheme for reachability with evaluation time $O(h)$*

Especially, if a class of planar digraphs have such an s-t-decompositions of constant height, they have an $O(\log n)$ bit labelling scheme for reachability.

3.1 Constructing an s-t-decomposition

The s-t-decomposition recursively chooses a face f and consequently a subgraph $H = bc(f)$ of the graph G induced by all vertices that can reach a vertex on f. Since G was embedded in the plane, the subgraph H is embedded in the plane, and all vertices of $G \setminus H$ lie in a unique face of H. We may choose a tree-cotree composition wisely, such that for each face of H, the restriction of T^* to the subfaces of that face is again a dual spanning tree (Lemma 3.10).

We also have to choose H carefully to ensure logarithmic height, and a limited alternation number on the frames. To ensure at most logarithmic height, we show two cases: 2-frame-nodes have only small children, while for 4-frame-nodes, we only need to ensure that their 4-frame children themselves are small.

Lemma 3.9. Let $G = (V, E)$ be a plane graph, let $G^* = (V^*, E^*)$ be its dual, let (T, T^*) be a tree/cotree decomposition of G, and let S be a subgraph of G such that $S \cap T$ is connected. Then the faces of S correspond to connected components of $T^* \setminus E^*[S]$.

Proof. Let S^* be the dual of S, then $S^* = G^*/(G^* \setminus E^*[S])$ and the claim is equivalent to saying that the components of $G^* \setminus E^*[S]$ correspond to the components of $T^* \setminus E^*[S]$. Consider a pair of faces $f_1, f_2 \in V^*$. Clearly, if they are in separate components of $G^* \setminus E^*[S]$, they are also in separate components in $T^* \setminus E^*[S]$. On the other hand, suppose f_1 and f_2 are in different components in $T^* \setminus E^*[S]$. Then there exists an edge $e^* \in E^*[S] \cap T^*$ separating them. The corresponding edge $e \in E[S]$ induces a cycle in T, which is also part of S since $S \cap T$ is connected. The dual to that cycle is an edge cut in G^* that separates f_1 from f_2. □

Lemma 3.10. Let T be a spanning tree where all edges point away from the source s of G, then for any node x in an st-decomposition of G, the subgraph T_x^* of T^* induced by $V^*[G_x^*]$ is a connected subtree of T^*.

Proof. If x is the root, this trivially holds. If x has a parent y, G_x^* corresponds to a face in S_y. Now $S_y \cap T$ is connected since S_y is the union of backward-closed graphs, and the result follows from Lemma 3.9. □

Lemma 3.11. Let x be a node in an st-decomposition whose parent frame F_x has alternation number 2, and let A^* be the set of faces in T_x^* incident to the target corner of F_x. Then for any child y of x:

$$A^* \subseteq V^*[T_y^*] \implies F_y \text{ has alternation number } 4.$$
$$A^* \not\subseteq V^*[T_y^*] \implies F_y \text{ has alternation number } 2.$$

Proof. Let t_x be the target corner of F_x and let A^* be the set of faces in ${T_x}^*$ incident to t_x. For any child y if x, F_y consists of a (possibly empty) segment of F_x and two directed paths that meet at a new target corner t_y. Each target corner of F_y must therefore be at either t_x or t_y. Now if $A^* \subseteq V^*[T_y^*]$, then both t_x and t_y are target corners of F_y, otherwise only t_y is. Either way the result follows. □

Lemma 3.12. Let x be a node in an st-decomposition whose parent frame F_x has alternation number 4, and let A^{0^*} and A^{1^*} be the sets of faces in T_x^* incident to the target corners of F_x. Then for any child y of x:

$$A^{0^*} \not\subseteq V^*[T_y^*] \vee A^{1^*} \not\subseteq V^*[T_y^*] \implies F_y \text{ has alternation number at most } 4.$$
$$A^{0^*} \not\subseteq V^*[T_y^*] \wedge A^{1^*} \not\subseteq V^*[T_y^*] \implies F_y \text{ has alternation number } 2.$$

Proof. Let t_x^0 and t_x^1 be the two target corners of F_x and for $i \in \{0, 1\}$ let A^{i^*} be the set of faces in ${T_x}^*$ incident to t_x^i. For any child y of x, F_y consists of a (possibly empty) segment of F_x and two directed paths that meet at a new target corner t_y. Each target corner of F_y must therefore be at either t_y, t_x^0, or t_x^1. Now if $A^{i^*} \not\subseteq V^*[T_y^*]$ for some $i \in \{0, 1\}$, then t_x^i is not a target corner of F_y. So the number of target corners in F_y is at least 1, and at most 3 minus the number of such i, and the result follows. □

proof of theorem 3.5. Let s be the source of G and let (T, T^*) be a tree/cotree decomposition of G such that all edges in T point away from s. The st-decomposition can be constructed recursively as follows. Start with the root. In each step we have a node x and by Lemma 3.10 the subgraph T_x^* induced in T^* by $V^*[G_x^*]$ is a tree. The goal is to select a face f_x such that for each child y:

- The alternation number of F_y is at most 4, and
- For each child z of y (and thus grandchild of x), $|T_z^*| \leq \frac{1}{2}|T_x^*|$.

If we can do this for all x, we are done. There are 3 cases:

x is the root Let f_x be the median of $T_x^* = T^*$. Then for each child y, $|T_y^*| \leq \frac{1}{2}|T_x^*|$, and, since $S_x = \mathrm{bc}(f_x)$ is a truncated s-t-graph with a single source, f_y has alternation number 2.

F_x has alternation number 2 Let f_x be the median of T_x^*. Then for each child y, $|T_y^*| \leq \frac{1}{2}|T_x^*|$, and, by Lemma 3.11, f_y has alternation number at most 4.

F_x **has alternation number** 4 Let t_0 and t_1 be the local targets of F_x and let $f_0, f_1 \in V^*[T_x^*]$ be (not necessarily distinct) faces incident to t_0 and t_1 respectively. Now choose f_x as the projection of the median m of T_x^* on the path $f_0, \ldots, f_1$ in T_x^*. By Lemma 3.12 this means that for any child y of x, the alternation number of the parent frame F_y is at most 4.
- If $f_x = m$ then $|T_y^*| \leq \frac{1}{2}|T_x^*|$.
- If $f_x \neq m$ and T_y^* is not the component of m in $T_x^* \setminus E^*[\mathrm{bc}(f_x)]$, then $|T_y^*| \leq \frac{1}{2}|T_x^*|$.
- If $f_x \neq m$, and T_y^* is the component of m in $T_x^* \setminus E^*[\mathrm{bc}(f_x)]$, then T_y^* contains neither f_0 nor f_1, so by Lemma 3.12 the parent frame F_y has alternation number at most 2 and we have just shown this means any child z of y has $|T_z^*| \leq \frac{1}{2}|T_y^*| \leq \frac{1}{2}|T_x^*|$. □

Note that this construction can be implemented in linear time by using ideas similar to [2].

3.2 Constructing a good s-t-decomposition in linear time

In the construction of an s-t-decomposition, a face is chosen, some edges are deleted, and new connected components of the dual graph arise. We then recurse on the new connected components of the dual graph. By Lemma 3.10 we can choose a tree/cotree-decomposition such that each component that arises is spanned by a subtree of the cotree.

To obtain linear construction time, we use a variation of the decremental tree connectivity algorithm from [2] to keep track of the subtrees of the cotree, and associate some information with each subtree. In particular, when T_x^* is a component at some point, we can in constant time find the node x.

For each node x we keep the set of target vertices on F_x (or $\emptyset$ if x is the root), and a face in T_x^* incident to each target in the set.

Build a top tree (see [1]) of height $O(\log n)$ over T^*, and let v_{n-i}^* be the i'th face that stops being boundary during the construction. Using this enumeration, the boundary faces of a cluster will be visited before boundary faces of their descendants. We use this ordering to find the splitting faces of the s-t-decomposition.

For each v_i^*, we can use the connectivity structure to find the relevant node x to split. We then need to choose the target face f_x defining the split. If x is the root or F_x is a 2-frame, we just set $f_x = v_i^*$. If F_x is a 4-frame, the information in x contains a pair of faces f_1, f_2 and we use a static nearest common ancestor data structure from Harel and Tarjan [8] to find the projection $f_x = \pi(v_i^*)$ of v_i^* on $f_1, \cdots, f_2$. Note that the projection of v_i^* is always contained in the same connected component as f_1, f_2, and thus, the data structure for the whole tree suffices to answer this query for the particular subtree.

Once f_x has been selected, we traverse the graph backwards from the vertices of f_x until we have found all the edges with destination in C_x. This search takes $|C_x|$ time. We delete these edges from the forest as we go along. Once we are done, we take all targets in C_x and select an incident face for each component it is incident to. This again takes $|C_x|$ time. If $f_x \neq v_i^*$ we try with v_i^* again, otherwise we move on to v_{i+1}^*.

Lemma 3.13. The s-t-decomposition constructed via the approach sketched above has no frame of alternation number > 4.

Proof. Components with 2-frames always have children with 2- and 4-frames. For components with 4-frames, this follows directly from Lemma 3.11, since we chose a splitting face on the cotree path between faces near the two targets. □

Lemma 3.14. The s-t-decomposition constructed via the approach sketched above has height $O(\log n)$.

Proof. Since the top-tree has height $O(\log n)$, choosing the boundary face v_i^* as a splitting face every time would result in a tree of the same height; $O(\log n)$. However, for each 4-frame, we might choose a face

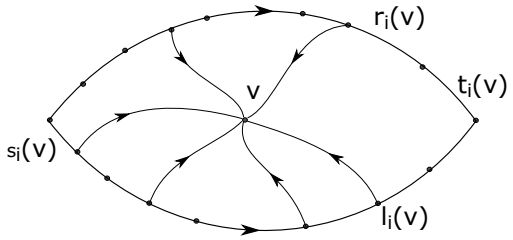

Figure 2: The best two vertices that can reach v on level i.

$f_x \neq v_i^*$ which is the projection of v_i^* on $f_1 \dots f_2$. As noted in Lemma 3.11, when this happens, v_i^* will lie in a child which has a 2-frame. But then, v_i^* will be the splitting face for that child. We thus increase the height by no more than a factor 2, and the s-t-decomposition has height $2O(\log n) = O(\log n)$. □

Lemma 3.15. Let $G = (V, E)$ be a plane single-source graph with source s, then we can construct a good s-t-decomposition of G in linear time.

Proof. Since the top-tree can be constructed in linear time, and since the decremental connectivity for trees takes linear time, and since the static nearest common ancestor data stucture is constructed in linear time and answers queries in constant time, the construction takes linear time. By Lemma 3.13 and 3.14, the resulting s-t-decomposition is good. □

3.3 2-frames

Definition 3.16. Let $\mathcal{T}$ be an st-decomposition of $G = (V, E)$. Then we can define a 2-*frame-decomposition* $\mathcal{T}_2$ by contracting each edge in $\mathcal{T}$ that corresponds to a 4-frame. For each node x in $\mathcal{T}_2$ that is contracted from a set of nodes $Y \subseteq \mathcal{T}$ define $C_x := \bigcup_{y \in Y} C_y$ and if x is not the root, define $F_x := F_{\text{lca}(Y)}$ and $E_x := E_{\text{lca}(Y)}$. Then F_x is a 2-frame, and we can define s_x to be the source corner, and t_x to be the target corner on F_x.

Definition 3.17. Let $(\mathcal{L}, \mathcal{R})$ be the partition of $\cup_{x \in \mathcal{T}_2} E_x$ defined as follows: For each $(u, v) \in \cup_{x \in \mathcal{T}_2} E_x$ let y be the node (if it exists) closest to the root of $\mathcal{T}_2$ such that $(u, v) \in E_y$ but u is not the target vertex of F_y. If y exists and (u, v) is incident to a corner on the clockwise disegment of F_y between s_y and t_y assign (u, v) to $\mathcal{R}$, otherwise assign (u, v) to $\mathcal{L}$.

Definition 3.18. Let $\mathcal{T}_2$ be an 2-frame-decomposition of $G = (V, E)$. For any vertex $v \in V$ define:

$c_2[v] :=$ The node x in $\mathcal{T}_2$ such that $v \in V[C_x]$
$d_2[v] :=$ The depth of $c_2[v]$ in $\mathcal{T}_2$

Definition 3.19. For any $0 \le i < d_2[v]$, let x be the ancestor of $c_2[v]$ at depth $i+1$ and define:

$$
\begin{aligned}
E_i(v) &:= E_x \\
L_i(v) &:= E_x \cap \mathcal{L} \\
R_i(v) &:= E_x \cap \mathcal{R} \\
\widehat{L}_i(v) &:= \left\{(w,w') \in L_i(v) \middle|\ w' \rightsquigarrow v\right\} \\
\widehat{R}_i(v) &:= \left\{(w,w') \in R_i(v) \middle|\ w' \rightsquigarrow v\right\} \\
\widehat{F}_i(v) &:= \widehat{L}_i(v) \cup \widehat{R}_i(v) \\
l_i(v) &:= \begin{cases} \bot & \text{if } \widehat{L}_i(v) = \emptyset \\ \text{the last vertex in } \operatorname{init}(\widehat{L}_i(v)) \text{ on the counterclockwise dipath of } F_x & \text{otherwise} \end{cases} \\
r_i(v) &:= \begin{cases} \bot & \text{if } \widehat{R}_i(v) = \emptyset \\ \text{the last vertex in } \operatorname{init}(\widehat{R}_i(v)) \text{ on the clockwise dipath of } F_x & \text{otherwise} \end{cases} \\
s_i(v) &:= \text{The vertex associated with } s_x \\
t_i(v) &:= \text{The vertex associated with } t_x
\end{aligned}
$$

Additionally, let $L_i(v)$ and $\widehat{L}_i(v)$ be totally ordered by the position of the starting vertices on the counterclockwise disegment of F_x and the clockwise order around each starting vertex. Similarly let $R_i(v)$ and $\widehat{R}_i(v)$ be totally ordered by the position of the starting vertices on the clockwise disegment of F_x and the counterclockwise order around each starting vertex.

The goal in this section is a data structure for efficiently computing $l_i(v)$ and $r_i(v)$ for $0 \le i < d_2[v]$.

Lemma 3.20. For any vertex $v \in V$ and $0 \le i < d_2[v]$: $\widehat{F}_i(v) \neq \emptyset$

Proof. Let x be the ancestor of $c_2[v]$ at depth $i+1$. Since G is a single-source graph, there is a path from s to v. This path must contain a vertex in $V[F_x]$. But then the edge following the last such vertex on the path must be in $\widehat{L}_i(v) \cup \widehat{R}_i(v)$ which is therefore nonempty. □

Lemma 3.21. For any $u, v \in V$ and $0 \le i < d_2[u]$: If $u \rightsquigarrow v$ then $\widehat{L}_i(u) \subseteq \widehat{L}_i(v)$ and $\widehat{R}_i(u) \subseteq \widehat{R}_i(v)$.

Proof. Since $u \rightsquigarrow v$, $c_2[u]$ is ancestor to $c_2[v]$ and so $L_i(u) = L_i(v)$ and hence $\widehat{L}_i(u) \subseteq \widehat{L}_i(v)$. Similarly, $R_i(u) = R_i(v)$ and $\widehat{R}_i(u) \subseteq \widehat{R}_i(v)$. □

Lemma 3.22. Given any vertex $v \in V$, $0 \le i < d_2[v]$, and $(w,w') \in E_i(v)$. Then:

$$
\begin{aligned}
(w,w') \in \widehat{L}_i(v) \quad &\implies \quad (w,w') \in \widehat{L}_{i'}(v) \text{ for all } i', d_2[w] \le i' < \min\left\{d_2[w'], d_2[v]\right\} \\
(w,w') \in \widehat{R}_i(v) \quad &\implies \quad (w,w') \in \widehat{R}_{i'}(v) \text{ for all } i', d_2[w] \le i' < \min\left\{d_2[w'], d_2[v]\right\}
\end{aligned}
$$

Proof. Let $j = d_2[w]$ and $k = \min\{d_2[w'], d_2[v]\}$. Clearly $(w,w') \in E_{i'}$ for all $j \le i' < k$. Suppose $(w,w') \in \widehat{L}_i(v) \subseteq L_i(v)$, then since $j \le i < k$ the definition give us $(w,w') \in L_{i'}(v)$ for all $j \le i' < k$. And since $w' \rightsquigarrow v$ this implies $(w,w') \in \widehat{L}_{i'}(v)$ for all $j \le i' < k$ and the result follows. The case for R is symmetric. □

Definition 3.23. For any vertex $v \in V$ let

$$p_l[v] := \begin{cases} \bot & \text{if } d_2[v] = 0 \\ l_{d_2[v]-1}(v) & \text{otherwise} \end{cases}$$
$$p_r[v] := \begin{cases} \bot & \text{if } d_2[v] = 0 \\ r_{d_2[v]-1}(v) & \text{otherwise} \end{cases}$$

and let T_l and T_r denote the rooted forests over V whose parent pointers are p_l and p_r respectively.

Definition 3.24. For any $v \in V \cup \{\bot\}$, and $i \geq 0$ let

$$l'_i(v) := \begin{cases} v & \text{if } v = \bot \vee d_2[v] \leq i \\ l'_i(p_l[v]) & \text{otherwise} \end{cases}$$
$$r'_i(v) := \begin{cases} v & \text{if } v = \bot \vee d_2[v] \leq i \\ r'_i(p_r[v]) & \text{otherwise} \end{cases}$$

Lemma 3.25. Let $v \in V$, and $i \geq 0$ be given, then

$$\begin{array}{lllll} i = d_2[v] - 1 & \implies & l'_i(v) = l_i(v) & \wedge & r'_i(v) = r_i(v) \\ i \leq d_2[v] - 1 & \implies & l'_i(v) \in \text{init}(\widehat{L}_i(v)) \cup \{\bot\} & \wedge & r'_i(v) \in \text{init}(\widehat{R}_i(v)) \cup \{\bot\} \\ i > d_2[v] - 1 & \implies & l'_i(v) = v & \wedge & r'_i(v) = v \end{array}$$

Proof. We will show this for l' only, as r' is completely symmetrical. If $i > d_2[v] - 1$ then $d_2[v] \leq i$ and we get $l'_i(v) = v$ directly from the definition of l'. Similarly if $i = d_2[v] - 1$ then $l'_i(v) = l'_i(p_l[v]) = l'_i(l_{d_2[v]-1}(v)) = l'_i(l_i(v)) = l_i(v) \in \text{init}(\widehat{L}_i(v)) \cup \{\bot\}$. Finally suppose $i < d_2[v] - 1$. If $l'_i(v) = \bot$ we are done, so suppose that is not the case. Let u be the child of $l'_i(v)$ in T_l that is ancestor to v. Then $l'_i(v) = l'_i(u) = p_l[u] = l_{d_2[u]-1}(u)$. By definition of $l_{d_2[u]-1}(u)$ there exists an edge $(w, w') \in \widehat{L}_{d_2[u]-1}$ where $w = l_{d_2[u]-1}(u)$ and $d_2[w] \leq i < d_2[w'] \leq d_2[u]$ and by setting $(v, i, (w, w')) = (u, d_2[u] - 1, (w, w'))$ in lemma 3.22 we get $(w, w') \in \widehat{L}_i(u)$, and therefore $l'_i(v) \in \text{init}(\widehat{L}_i(u))$. But since $u \rightsquigarrow v$ we have $\widehat{L}_i(u) \subseteq \widehat{L}_i(v)$ by Lemma 3.21 and we are done. □

Lemma 3.26. Let $v \in V$ and $0 \leq i \leq j$ then

$$l'_i(l'_j(v)) = l'_i(v) \qquad \wedge \qquad r'_i(r'_j(v)) = r'_i(v)$$

Proof. $l'_j(v)$ is on the path from v to $l'_i(v)$ in T_l, so this follows trivially from the recursion. The case for r' is symmetric. □

Lemma 3.27. Let $v \in V$, and $0 \leq i < d_2[v] - 1$, then

$$l_i(v) = \bot \implies l'_i(l_{i+1}(v)) = \bot \qquad \wedge \qquad r_i(v) = \bot \implies r'_i(r_{i+1}(v)) = \bot$$

Proof. If $l_i(v) = \bot$ then $\widehat{L}_i(v) = \emptyset$, so either $l_{i+1}(v) = \bot$ implying $l'_i(l_{i+1}(v)) = \bot$ by the definition of l', or $l_{i+1}(v) \notin \text{init}(\widehat{L}_i(v))$ so $d_2[l_{i+1}(v)] = i + 1$ and by Lemma 3.25 and Lemma 3.21 $l'_i(l_{i+1}(v)) \in \text{init}(\widehat{L}_i(l_{i+1}(v))) \cup \{\bot\} \subseteq \text{init}(\widehat{L}_i(v)) \cup \{\bot\} = \{\bot\}$ so again $l'_i(l_{i+1}(v)) = \bot$. The case for r is symmetric. □

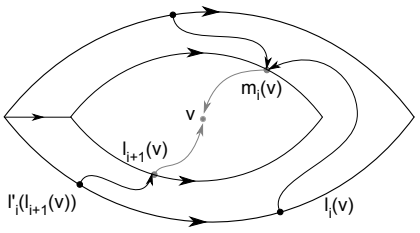

Figure 3: The best path from $L_i(v)$ goes via $r_{i+1}(v)$.

Lemma 3.28 (Crossing lemma)**.** Let $v \in V$, and $0 \leq i < d_2[v] - 1$.

$$\begin{aligned} l_i(v) \neq l'_i(l_{i+1}(v)) \implies & \quad l_i(v) = l'_i(m) \wedge r_i(v) = r'_i(m) \wedge d_2[m] = i+1 \\ & \text{where } m = r_{i+1}(v) \neq \bot \\ r_i(v) \neq r'_i(r_{i+1}(v)) \implies & \quad l_i(v) = l'_i(m) \wedge r_i(v) = r'_i(m) \wedge d_2[m] = i+1 \\ & \text{where } m = l_{i+1}(v) \neq \bot \end{aligned}$$

Proof. Suppose $l_i(v) \neq l'_i(l_{i+1}(v))$ (the case $r_i(v) \neq r'_i(r_{i+1}(v))$ is symmetrical). Then $l_i(v) \neq \bot$ by lemma 3.27. Thus there is a last edge $(w, w') \in \widehat{L}_i(v)$ with $w = l_i(v)$ and $d_2[w] \leq i < d_2[w']$ and a path $P = w' \rightsquigarrow v$.

Now $(w, w') \notin E_{i+1}(v)$ since otherwise by Definition 3.19 $(w, w') \in L_{i+1}(v)$ and since $w' \rightsquigarrow v$ even $(w, w') \in \widehat{L}_{i+1}(v)$ implying $l_i(v) = l_{i+1}(v)$ and thus $l_i(v) = l'_i(l_{i+1}(v))$ by lemma 3.25, contradicting our assumption.

Since $(w, w') \notin E_{i+1}(v)$, the path P must cross $\widehat{F}_{i+1}(v)$. Let (u, u') be the last edge in $P \cap \widehat{F}_{i+1}(v)$. Then $w' \rightsquigarrow u$ so $d_2[u] \geq i+1$ and $(u, u') \notin L_{i+1}(v)$ since otherwise $d_2[l_{i+1}(v)] = i+1$ and hence by Lemma 3.25 $l_i(v) = l'_i(l_{i+1}(v))$, again contradicting our assumption. Since $\widehat{F}_{i+1}(v) \neq \emptyset$, we therefore have $(u, u') \in \widehat{R}_{i+1}(v)$.

But then we can choose P so it goes through (m, m') where $m = r_{i+1}(v) \neq \bot$. Now $i+1 \leq d_2[w'] \leq d_2[r_{i+1}(v)] \leq i+1$ so $d_2[m] = i+1$.

Let e be the last edge in $\widehat{R}_i(v)$ then any path $r_i(v) \rightsquigarrow v$ that starts with e crosses $P \cup \widehat{R}_{i+1}(v)$, implying that there exists such a path that contains (m, m') and thus $r_i(v) = r_i(m)$. Since $d_2[m] = i+1$, then $l_i(v) = l'_i(m)$ and $r_i(v) = r'_i(m)$ follows from lemma 3.25. □

Definition 3.29. Let $v \in V$ and $0 \leq i < d_2[v]$.

$$m_i(v) := \begin{cases} v & \text{if } i+1 = d_2[v] \\ l_{i+1}(v) & \text{if } i+1 < d_2[v] \wedge r_i(v) \neq r'_i(r_{i+1}(v)) \\ r_{i+1}(v) & \text{if } i+1 < d_2[v] \wedge l_i(v) \neq l'_i(l_{i+1}(v)) \\ m_{i+1}(v) & \text{otherwise} \end{cases}$$

Corollary 3.30. Let $v \in V$ and $0 \leq i < d_2[v] - 1$. If $l_i(v) \neq l'_i(l_{i+1}(v))$ or $r_i(v) \neq r'_i(r_{i+1}(v))$ then

$$l_i(v) = l'_i(m_i(v)) \qquad \wedge \qquad r_i(v) = r'_i(m_i(v)) \qquad \wedge \qquad d_2[m_i(v)] = i+1$$

Proof. This is just a reformulation of lemma 3.28 in terms of $m_i(v)$. □

Lemma 3.31. For any vertex $v \in V$ and $0 \leq i < d_2[v]$

$$l_i(v) = l'_i(m_i(v)) \qquad \wedge \qquad r_i(v) = r'_i(m_i(v))$$

Proof. The proof is by induction on j, the number of times the "otherwise" case is used before reaching one of the other cases when expanding the recursive definition of $m_i(v)$.

For $j = 0$, either $i + 1 = d_2[v]$ and the result follows from Lemma 3.25, or $i + 1 < d_2[v]$ and $l_i(v) \neq l'_i(l_{i+1}(v))$ or $r_i(v) \neq r'_i(r_{i+1}(v))$. In either case we have by Corollary 3.30, that $l_i(v) = l'_i(m_i(v))$ and $r_i(v) = r'_i(m_i(v))$.

For $j > 0$ we have $i + 1 < d_2[v]$ and $l_i(v) = l'_i(l_{i+1}(v))$ and $r_i(v) = r'_i(r_{i+1}(v))$ and $m_i(v) = m_{i+1}(v)$. By induction we can assume that $l_{i+1}(v) = l'_{i+1}(m_{i+1}(v))$ and $r_{i+1}(v) = r'_{i+1}(m_{i+1}(v))$. Then by Lemma 3.26, $l'_i(l_{i+1}(v)) = l'_i(l'_{i+1}(m_{i+1}(v))) = l'_i(m_{i+1}(v)) = l'_i(m_i(v))$, showing that $l_i(v) = l'_i(m_i(v))$ as desired. The case for r is symmetric. □

Definition 3.32. For any vertex $v \in V$, let

$$M[v] := \{i |\ 0 < i < d_2[v] \wedge m_{i-1}(v) \neq m_i(v)\}$$
$$p_m[v] := \begin{cases} \perp & \text{if } M[v] = \emptyset \\ m_{\max M[v]-1}(v) & \text{otherwise} \end{cases}$$

And define T_m as the rooted forest over V whose parent pointers are p_m.

Theorem 3.33. *There exists a practical RAM data structure that for any good st-decomposition of a graph with n vertices uses $\mathcal{O}(n)$ words of $\mathcal{O}(\log n)$ bits and can answer $l_i(v)$ and $r_i(v)$ queries in constant time.*

Proof. For any vertex $v \in V$, let

$$D_l[v] := \{i |\ v \text{ has a proper ancestor } w \text{ in } T_l \text{ with } d_2[w] = i\}$$
$$D_r[v] := \{i |\ v \text{ has a proper ancestor } w \text{ in } T_r \text{ with } d_2[w] = i\}$$

Now, store levelancestor structures for each of T_l, T_r, and T_m, together with $d_2[v]$, $D_l[v]$, $D_r[v]$, and $M[v]$ for each vertex. Since the height of the st-decomposition is $\mathcal{O}(\log n)$ each of $D_l[v]$, $D_r[v]$, and $M[v]$ can be represented in a single $\mathcal{O}(\log n)$-bit word.

This representation allows us to find $d_2[m_i(v)] = \text{succ}(M[v] \cup \{d_2[v]\}, i)$ in constant time, as well as computing the depth in T_m of $m_i(v)$. Then using the levelancestor structure for T_m we can compute $m_i(v)$ in constant time.

Similarly, this representation of the $D_l[v]$ set lets us compute the depth in T_l of $l'_i(v)$ in constant time, and with the levelancestor structure that lets us compute $l'_i(v)$ in constant time. A symmetric argument shows that we can compute $r'_i(v)$ in constant time.

Finally, lemma 3.31 says we can compute $l_i(v)$ and $r_i(v)$ in constant time given constant-time functions for l', r', and m. □

3.4 4-frames

Definition 3.34. Let x be a node in an s-t-decomposition such that F_x is a 4-frame, and let y be its parent. Let s^0_x and s^1_x be the source corners on F_x and let t^0_x and t^1_x be the target corners on F_x, numbered such that their clockwise cyclic order on F_x is $s^0_x, t^0_x, s^1_x, t^1_x$, and such that if F_y is a 4-frame there is an $\alpha \in \{0, 1\}$ so $t^\alpha_x = t^\alpha_y$.

Definition 3.35. Let $\mathcal{E}_4$ be the set of edges (u, v) such if x is the node in the s-t-decomposition that contains v, then $(u, v) \in E_x$ and F_x is a 4-frame. Let $(\mathcal{L}^0, \mathcal{R}^0, \mathcal{L}^1, \mathcal{R}^1)$ be the partition of $\mathcal{E}_4$ defined as follows: For each $(u, v) \in \mathcal{E}_4$ let x be the node such that $v \in C_x$, and let y be the node (if it exists) closest to the root of $\mathcal{T}$ such that

- For any z that is ancestor to x and descendent to y, F_z is a 4-frame.
- $(u, v) \in E_y$.
- u is not a target vertex of F_y.

If y exists, then (u, v) is incident to a corner c on F_y. If there is an $\alpha \in \{0, 1\}$ such that c is on the clockwise disegment of F_y between s_y^α and t_y^α we assign (u, v) to $\mathcal{R}^\alpha$. Otherwise there must be an $\alpha \in \{0, 1\}$ such that c is on the counterclockwise disegment of F_y between $s_y^{1-\alpha}$ and t_y^α, and we assign (u, v) to $\mathcal{L}^\alpha$. If no such y exists, (u, v) must be incident to t_x^α for some $\alpha \in \{0, 1\}$ and we (arbitrarily) assign (u, v) to $\mathcal{L}^\alpha$.

Definition 3.36. Let $\mathcal{T}$ be an st-decomposition of $G = (V, E)$. For any vertex $v \in V$ define:

$$
\begin{aligned}
c[v] &:= \text{The node } x \text{ in } \mathcal{T} \text{ such that } v \in V[C_x] \\
d[v] &:= \text{The depth of } c[v] \text{ in } \mathcal{T} \\
J_2[v] &:= \{\text{depth}(x) |\ x \text{ is a non-root ancestor to } c[v] \text{ in } \mathcal{T} \text{ and } F_x \text{ is a 2-frame}\} \\
j_2[v] &:= \max(J_2[v]))
\end{aligned}
$$

The number $j_2[v]$ is especially useful for 4-frame nodes. On the path from the root to the component of v in the s-t-decomposition tree, there will be a last component whose frame is a 2-frame. We call the depth of the next component on the path $j_2[v]$. If $c[v]$ has a 4-frame, then for the rest of the path, that is, depth i with $j_2[v] \le i < d[v]$, we will have 4-frames nested in 4-frames, which gives a lot of useful structure.

Definition 3.37. For any $j_2[v] \le i < d[v]$ and $\alpha \in \{0, 1\}$, let x be the ancestor of $c[v]$ at depth $i + 1$ and define:

$$
\begin{aligned}
E_i(v) &:= E_x \\
L_i^\alpha(v) &:= E_x \cap \mathcal{L}^\alpha \\
R_i^\alpha(v) &:= E_x \cap \mathcal{R}^\alpha \\
\widehat{L}_i^\alpha(v) &:= \left\{(w, w') \in L_i^\alpha(v) \middle|\ w' \rightsquigarrow v\right\} \\
\widehat{R}_i^\alpha(v) &:= \left\{(w, w') \in R_i^\alpha(v) \middle|\ w' \rightsquigarrow v\right\} \\
\widehat{F}_i(v) &:= \widehat{L}_i^0(v) \cup \widehat{R}_i^0(v) \cup \widehat{L}_i^1(v) \cup \widehat{R}_i^1(v) \\
l_i^\alpha(v) &:= \begin{cases} \bot & \text{if } \widehat{L}_i^\alpha(v) = \emptyset \\ \text{the last vertex in } \text{init}(\widehat{L}_i^\alpha(v)) \text{ on the counterclockwise dipath of } F_x & \text{otherwise} \end{cases} \\
r_i^\alpha(v) &:= \begin{cases} \bot & \text{if } \widehat{R}_i^\alpha(v) = \emptyset \\ \text{the last vertex in } \text{init}(\widehat{R}_i^\alpha(v)) \text{ on the clockwise dipath of } F_x & \text{otherwise} \end{cases} \\
s_i^\alpha(v) &:= \text{The vertex associated with } s_x^\alpha \\
t_i^\alpha(v) &:= \text{The vertex associated with } t_x^\alpha
\end{aligned}
$$

Additionally, let $L_i^\alpha(v)$ and $\widehat{L}_i^\alpha(v)$ be totally ordered by the position of the starting vertices on the counterclockwise disegment of F_x and the clockwise order around each starting vertex. Similarly let $R_i^\alpha(v)$ and $\widehat{R}_i^\alpha(v)$ be totally ordered by the position of the starting vertices on the clockwise disegment of F_x and the counterclockwise order around each starting vertex.

We know from Section 3.3 that we can find the relevant vertices on each 2-frame surrounding v. The goal in this section is a data structure for efficiently computing $l_i^\alpha(v)$ and $r_i^\alpha(v)$ for $j_2[v] \le i < d[v]$.

Lemma 3.38. For any vertex $v \in V$ and $j_2[v] \leq i < d[v]$: $\widehat{F}_i(v) \neq \emptyset$

Proof. Let x be the ancestor of $c[v]$ at depth $i+1$. Since G is a single-source graph, there is a path from s to v. This path must contain a vertex in $V[F_x]$. But then the edge following the last such vertex on the path must be in $\widehat{L}_i^0(v) \cup \widehat{R}_i^0(v) \cup \widehat{L}_i^1(v) \cup \widehat{R}_i^1(v)$ which is therefore nonempty. □

Lemma 3.39. For any $u, v \in V$, $j_2[v] \leq i < d[u]$, and $\alpha \in \{0,1\}$: If $u \rightsquigarrow v$ then $\widehat{L}_i^\alpha(u) \subseteq \widehat{L}_i^\alpha(v)$ and $\widehat{R}_i^\alpha(u) \subseteq \widehat{R}^a lpha_i(v)$.

Proof. Since $u \rightsquigarrow v$, $c[u]$ is ancestor to $c[v]$ and so $L_i^\alpha(u) = L_i^\alpha(v)$ and hence $\widehat{L}_i^\alpha(u) \subseteq \widehat{L}_i^\alpha(v)$. Similarly, $R_i^\alpha(u) = R_i^\alpha(v)$ and $\widehat{R}_i^\alpha(u) \subseteq \widehat{R}_i^\alpha(v)$. □

Lemma 3.40. Given any vertex $v \in V$, $j_2[v] \leq i < d[v]$, $\alpha \in \{0,1\}$, and $(w, w') \in E_i(v)$. Then:

$$(w,w') \in \widehat{L}_i^\alpha(v) \implies (w,w') \in \widehat{L}_{i'}^\alpha(v) \text{ for all } i', \max\{d[w], j_2[v]\} \leq i' < \min\left\{d[w'], d[v]\right\}$$
$$(w,w') \in \widehat{R}_i^\alpha(v) \implies (w,w') \in \widehat{R}_{i'}^\alpha(v) \text{ for all } i', \max\{d[w], j_2[v]\} \leq i' < \min\left\{d[w'], d[v]\right\}$$

Proof. Let $j = \max\{d[w], j_2[v]\}$ and $k = \min\{d[w'], d[v]\}$. Clearly $(w,w') \in E_{i'}$ for all $j \leq i' < k$. Suppose $(w,w') \in \widehat{L}_i^\alpha(v) \subseteq L_i^\alpha(v)$, then since $j \leq i < k$ the definition give us $(w,w') \in L_{i'}^\alpha(v)$ for all $j \leq i' < k$. And since $w' \rightsquigarrow v$ this implies $(w,w') \in \widehat{L}_{i'}^\alpha(v)$ for all $j \leq i' < k$ and the result follows. The case for R is symmetric. □

Definition 3.41. For any vertex $v \in V$ and $\alpha \in \{0,1\}$ let

$$p_l^\alpha[v] := \begin{cases} \bot & \text{if } d[v] = 0 \vee \ F_{d[v]-1}(v) \text{ is a 2-frame} \\ l^\alpha_{d[v]-1}(v) & \text{otherwise} \end{cases}$$
$$p_r^\alpha[v] := \begin{cases} \bot & \text{if } d[v] = 0 \vee \ F_{d[v]-1}(v) \text{ is a 2-frame} \\ r^\alpha_{d[v]-1}(v) & \text{otherwise} \end{cases}$$

and let T_l^α and T_r^α denote the rooted forests over V whose parent pointers are p_l^α and p_r^α respectively.

Definition 3.42. For any $v \in V \cup \{\bot\}$, $\alpha \in \{0,1\}$, and $i \geq j_2[v]$ let

$$l_i'^\alpha(v) := \begin{cases} v & \text{if } v = \bot \vee d[v] \leq i \\ l_i'^\alpha(p_l^\alpha[v]) & \text{otherwise} \end{cases}$$
$$r_i'^\alpha(v) := \begin{cases} v & \text{if } v = \bot \vee d[v] \leq i \\ r_i'^\alpha(p_r^\alpha[v]) & \text{otherwise} \end{cases}$$

Lemma 3.43. Let $v \in V$, $\alpha \in \{0,1\}$, and $i \geq j_2[v]$ be given, then

$$i = d[v]-1 \implies l_i'^\alpha(v) = l_i^\alpha(v) \ \wedge \ r_i'^\alpha(v) = r_i^\alpha(v)$$
$$i \leq d[v]-1 \implies l_i'^\alpha(v) \in \text{init}(\widehat{L}_i^\alpha(v)) \cup \{\bot\} \ \wedge \ r_i'^\alpha(v) \in \text{init}(\widehat{R}_i^\alpha(v)) \cup \{\bot\}$$
$$i > d[v]-1 \implies l_i'^\alpha(v) = v \ \wedge \ r_i'^\alpha(v) = v$$

Proof. We will show this for l' only, as r' is completely symmetrical. If $i > d[v] - 1$ then $d[v] \leq i$ and we get $l'^{\alpha}_i(v) = v$ directly from the definition of l'. Similarly if $i = d[v] - 1$ then $l'^{\alpha}_i(v) = l'^{\alpha}_i(p^{\alpha}_l[v]) = l'^{\alpha}_i(l^{\alpha}_{d[v]-1}(v)) = l'^{\alpha}_i(l^{\alpha}_i(v)) = l^{\alpha}_i(v) \in \text{init}(\widehat{L}^{\alpha}_i(v)) \cup \{\bot\}$. Finally suppose $i < d[v] - 1$. If $l'^{\alpha}_i(v) = \bot$ we are done, so suppose that is not the case. Let u be the child of $l'^{\alpha}_i(v)$ in T_l that is ancestor to v. Then $l'^{\alpha}_i(v) = l'^{\alpha}_i(u) = p^{\alpha}_l[u] = l^{\alpha}_{d[u]-1}(u)$. By definition of $l^{\alpha}_{d[u]-1}(u)$ there exists an edge $(w, w') \in \widehat{L}^{\alpha}_{d[u]-1}$ where $w = l^{\alpha}_{d[u]-1}(u)$ and $d[w] \leq i < d[w'] \leq d[u]$ and by setting $(v, i, (w, w')) = (u, d[u] - 1, (w, w'))$ in lemma 3.40 we get $(w, w') \in \widehat{L}^{\alpha}_i(u)$, and therefore $l'^{\alpha}_i(v) \in \text{init}(\widehat{L}^{\alpha}_i(u))$. But since $u \rightsquigarrow v$ we have $\widehat{L}^{\alpha}_i(u) \subseteq \widehat{L}^{\alpha}_i(v)$ by Lemma 3.39 and we are done. □

Lemma 3.44. Let $v \in V$, $\alpha \in \{0, 1\}$, and $j_2[v] \leq i \leq j$ then

$$l'^{\alpha}_i(l'^{\alpha}_j(v)) = l'^{\alpha}_i(v) \qquad \wedge \qquad r'^{\alpha}_i(r'^{\alpha}_j(v)) = r'^{\alpha}_i(v)$$

Proof. $l'^{\alpha}_j(v)$ is on the path from v to $l'^{\alpha}_i(v)$ in T_l, so this follows trivially from the recursion. The case for r' is symmetric. □

Lemma 3.45. Let $v \in V$, $\alpha \in \{0, 1\}$, and $j_2[v] \leq i < d[v] - 1$, then

$$l^{\alpha}_i(v) = \bot \implies l'^{\alpha}_i(l^{\alpha}_{i+1}(v)) = \bot \qquad \wedge \qquad r^{\alpha}_i(v) = \bot \implies r'^{\alpha}_i(r^{\alpha}_{i+1}(v)) = \bot$$

Proof. If $l^{\alpha}_i(v) = \bot$ then $\widehat{L}^{\alpha}_i(v) = \emptyset$, so either $l^{\alpha}_{i+1}(v) = \bot$ implying $l'^{\alpha}_i(l^{\alpha}_{i+1}(v)) = \bot$ by the definition of l', or $l^{\alpha}_{i+1}(v) \notin \text{init}(\widehat{L}^{\alpha}_i(v))$ so $d[l^{\alpha}_{i+1}(v)] = i + 1$ and by Lemma 3.43 $l'^{\alpha}_i(l^{\alpha}_{i+1}(v)) \in \text{init}(\widehat{L}^{\alpha}_i(l^{\alpha}_{i+1}(v))) \cup \{\bot\} \subseteq \text{init}(\widehat{L}^{\alpha}_i(v)) \cup \{\bot\} = \{\bot\}$ so again $l'^{\alpha}_i(l^{\alpha}_{i+1}(v)) = \bot$. The case for r is symmetric. □

Lemma 3.46 (Crossing lemma)**.** Let $v \in V$, $\alpha \in \{0, 1\}$, and $j_2[v] \leq i < d[v] - 1$.

$$\begin{aligned} l^{\alpha}_i(v) \neq l'^{\alpha}_i(l^{\alpha}_{i+1}(v)) &\implies l^{\alpha}_i(v) = l'^{\alpha}_i(m) \wedge r^{\alpha}_i(v) = r'^{\alpha}_i(m) \wedge d[m] = i + 1 \\ &\quad \text{where } m = r^{\alpha}_{i+1}(v) \neq \bot \\ r^{\alpha}_i(v) \neq r'^{\alpha}_i(r^{\alpha}_{i+1}(v)) &\implies l^{\alpha}_i(v) = l'^{\alpha}_i(m) \wedge r^{\alpha}_i(v) = r'^{\alpha}_i(m) \wedge d[m] = i + 1 \\ &\quad \text{where } m = l^{\alpha}_{i+1}(v) \neq \bot \end{aligned}$$

Proof. Suppose $l^{\alpha}_i(v) \neq l'^{\alpha}_i(l^{\alpha}_{i+1}(v))$ (the case $r^{\alpha}_i(v) \neq r'^{\alpha}_i(r^{\alpha}_{i+1}(v))$ is symmetrical). Then $l^{\alpha}_i(v) \neq \bot$ by lemma 3.45. Thus there is a last edge $(w, w') \in \widehat{L}^{\alpha}_i(v)$ with $w = l^{\alpha}_i(v)$ and $d[w] \leq i < d[w']$ and a path $P = w' \rightsquigarrow v$.

Now $(w, w') \notin E_{i+1}(v)$ since otherwise by Definition 3.37 $(w, w') \in L^{\alpha}_{i+1}(v)$ and since $w' \rightsquigarrow v$ even $(w, w') \in \widehat{L}^{\alpha}_{i+1}(v)$ implying $l^{\alpha}_i(v) = l^{\alpha}_{i+1}(v)$ and thus $l^{\alpha}_i(v) = l'^{\alpha}_i(l^{\alpha}_{i+1}(v))$ by lemma 3.43, contradicting our assumption.

Since $(w, w') \notin E_{i+1}(v)$, the path P must cross $\widehat{F}_{i+1}(v)$. Let (u, u') be the last edge in $P \cap \widehat{F}_{i+1}(v)$. Then $w' \rightsquigarrow u$ so $d[u] \geq i + 1$ and $(u, u') \notin L^{\alpha}_{i+1}(v)$ since otherwise $d[l^{\alpha}_{i+1}(v)] = i + 1$ and hence by Lemma 3.43 $l^{\alpha}_i(v) = l'^{\alpha}_i(l^{\alpha}_{i+1}(v))$, again contradicting our assumption.

Also, $t^{\alpha}_i(v) \neq t^{\alpha}_{i+1}(v)$ because $t^{\alpha}_i(v) = t^{\alpha}_{i+1}(v)$ would imply $(w, w') \in L^{\alpha}_{i+1}(v) \cup \{\bot\}$ which we have just shown is not the case.

Since $t^{\alpha}_i(v) \neq t^{\alpha}_{i+1}(v)$, then by definition $t^{1-\alpha}_i(v) = t^{1-\alpha}_{i+1}(v)$ and hence $L^{1-\alpha}_{i+1}(v) \subseteq L^{1-\alpha}_i(v)$ and $R^{1-\alpha}_{i+1}(v) \subseteq R^{1-\alpha}_i(v)$, implying $d[w''] \leq i$ for all $w'' \in L^{1-\alpha}_{i+1}(v) \cup R^{1-\alpha}_{i+1}(v)$. Thus, $(u, u') \notin L^{1-\alpha}_{i+1}(v) \cup R^{1-\alpha}_{i+1}(v)$ since $d[u] > i$, and we can conclude that $(u, u') \in \widehat{R}^{\alpha}_{i+1}(v)$.

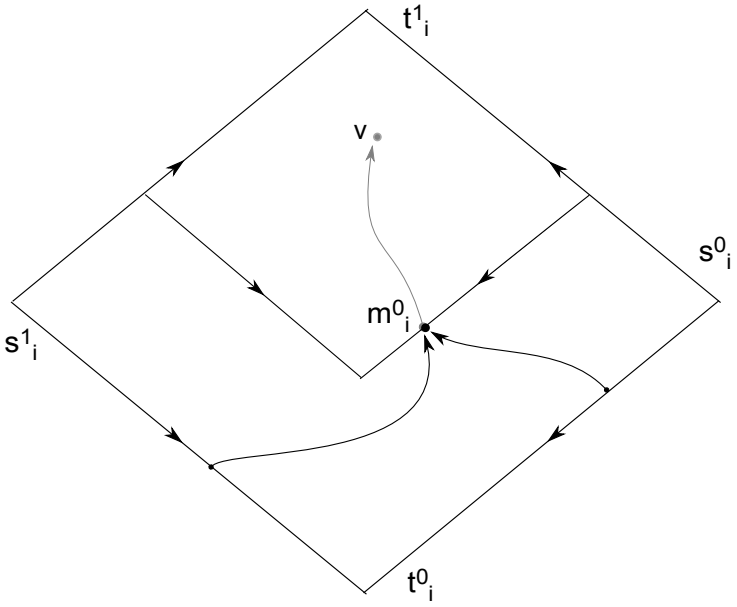

Figure 4: Sometimes the best path from $L_i^0(v)$ to v must go through $R_{i+1}^0(v)$.

But then we can choose P so it goes through (m, m') where $m = r_{i+1}^\alpha(v) \neq \perp$. Now $i+1 \leq d[w'] \leq d[r_{i+1}^\alpha(v)] \leq i+1$ so $d[m] = i+1$.

Let e be the last edge in $\widehat{R}_i^\alpha(v)$ then any path $r_i^\alpha(v) \rightsquigarrow v$ that starts with e crosses $P \cup \widehat{R}_{i+1}^\alpha(v)$, implying that there exists such a path that contains (m, m') and thus $r_i^\alpha(v) = r_i^\alpha(m)$. Since $d[m] = i+1$, then $l_i^\alpha(v) = l_i'^\alpha(m)$ and $r_i^\alpha(v) = r_i'^\alpha(m)$ follows from lemma 3.43. □

Definition 3.47. Let $v \in V$, $\alpha \in \{0,1\}$, and $0 \leq i < d[v]$.

$$m_i^\alpha(v) := \begin{cases} v & \text{if } i+1 = d[v] \\ l_{i+1}^\alpha(v) & \text{if } i+1 < d[v] \wedge r_i^\alpha(v) \neq r_i'^\alpha(r_{i+1}^\alpha(v)) \\ r_{i+1}^\alpha(v) & \text{if } i+1 < d[v] \wedge l_i^\alpha(v) \neq l_i'^\alpha(l_{i+1}^\alpha(v)) \\ m_{i+1}^\alpha(v) & \text{otherwise} \end{cases}$$

Corollary 3.48. Let $v \in V$, $\alpha \in \{0,1\}$, and $j_2[v] \leq i < d[v] - 1$. If $l_i^\alpha(v) \neq l_i'^\alpha(l_{i+1}^\alpha(v))$ or $r_i^\alpha(v) \neq r_i'^\alpha(r_{i+1}^\alpha(v))$ then

$$l_i^\alpha(v) = l_i'^\alpha(m_i^\alpha(v)) \quad \wedge \quad r_i^\alpha(v) = r_i'^\alpha(m_i^\alpha(v)) \quad \wedge \quad d[m_i^\alpha(v)] = i+1$$

Proof. This is just a reformulation of lemma 3.46 in terms of $m_i^\alpha(v)$. □

Lemma 3.49. For any vertex $v \in V$, $\alpha \in \{0,1\}$, and $j_2[v] \leq i < d[v]$

$$l_i^\alpha(v) = l_i'^\alpha(m_i^\alpha(v)) \quad \wedge \quad r_i^\alpha(v) = r_i'^\alpha(m_i^\alpha(v))$$

Proof. The proof is by induction on j, the number of times the "otherwise" case is used before reaching one of the other cases when expanding the recursive definition of $m_i(v)$.

For $j = 0$, either $i + 1 = d[v]$ and the result follows from Lemma 3.43, or $i + 1 < d[v]$ and $l_i(v) \neq l'_i(l_{i+1}(v))$ or $r_i(v) \neq r'_i(r_{i+1}(v))$. In either case we have by Corollary 3.48, that $l^\alpha_i(v) = l'^\alpha_i(m^\alpha_i(v))$ and $r^\alpha_i(v) = r'^\alpha_i(m^\alpha_i(v))$.

For $j > 0$ we have $i + 1 < d[v]$ and $l_i(v) = l'_i(l_{i+1}(v))$ and $r_i(v) = r'_i(r_{i+1}(v))$ and $m_i(v) = m_{i+1}(v)$. By induction we can assume that $l^\alpha_{i+1}(v) = l'^\alpha_{i+1}(m^\alpha_{i+1}(v))$ and $r^\alpha_{i+1}(v) = r'^\alpha_{i+1}(m^\alpha_{i+1}(v))$. Then by Lemma 3.44, $l'^\alpha_i(l^\alpha_{i+1}(v)) = l'^\alpha_i(l'^\alpha_{i+1}(m^\alpha_{i+1}(v))) = l'^\alpha_i(m^\alpha_{i+1}(v)) = l'^\alpha_i(m^\alpha_i(v))$, showing that $l^\alpha_i(v) = l'^\alpha_i(m^\alpha_i(v))$ as desired. The case for r is symmetric. □

Definition 3.50. For any vertex $v \in V$, and $\alpha \in \{0, 1\}$ let

$$M^\alpha[v] := \left\{ i \middle| \, j_2[v] < i < d[v] \wedge m^\alpha_{i-1}(v) \neq m^\alpha_i(v) \right\}$$

$$p^\alpha_m[v] := \begin{cases} \perp & \text{if } M^\alpha[v] = \emptyset \\ m^\alpha_{\max M^\alpha[v]-1}(v) & \text{otherwise} \end{cases}$$

And define T^α_m as the rooted forest over V whose parent pointers are p^α_m.

Theorem 3.51. *There exists a practical RAM data structure that for any good st-decomposition of a graph with n vertices uses $\mathcal{O}(n)$ words of $\mathcal{O}(\log n)$ bits and can answer $l^\alpha_i(v)$ and $r^\alpha_i(v)$ queries in constant time.*

Proof. For any vertex $v \in V$, and $\alpha \in \{0, 1\}$ let

$$D^\alpha_l[v] := \{i| \text{ } v \text{ has a proper ancestor } w \text{ in } T^\alpha_l \text{ with } d[w] = i\}$$
$$D^\alpha_r[v] := \{i| \text{ } v \text{ has a proper ancestor } w \text{ in } T^\alpha_r \text{ with } d[w] = i\}$$

Now, store levelancestor structures for each of T^α_l, T^α_r, and T^α_m, together with $d[v]$, $j_2[v]$, $J_2[v]$, $D^\alpha_l[v]$, $D^\alpha_r[v]$, and $M^\alpha[v]$ for each vertex. Since the height of the st-decomposition is $\mathcal{O}(\log n)$ each of $J_2[v]$, $D^\alpha_l[v]$, $D^\alpha_r[v]$, and $M^\alpha[v]$ can be represented in a single $\mathcal{O}(\log n)$-bit word.

This representation allows us to find $d[m^\alpha_i(v)] = \text{succ}(M^\alpha[v] \cup \{d[v]\}, i)$ in constant time, as well as computing the depth in T^α_m of $m^\alpha_i(v)$. Then using the levelancestor structure for T^α_m we can compute $m^\alpha_i(v)$ in constant time.

Similarly, this representation of the $D^\alpha_l[v]$ set lets us compute the depth in T^α_l of $l'^\alpha_i(v)$ in constant time, and with the levelancestor structure that lets us compute $l'^\alpha_i(v)$ in constant time. A symmetric argument shows that we can compute $r'^\alpha_i(v)$ in constant time.

Finally, lemma 3.49 says we can compute $l^\alpha_i(v)$ and $r^\alpha_i(v)$ in constant time given constant-time functions for l', r', and m. □

4 Acyclic planar In- out- graphs

For an in-out-graph G we have a source, s, that can reach all vertices of outdegree 0. Given such a source, s, we may assign all vertices a colour: A vertex is green if it can be reached from s, and red otherwise. We may also colour the directed edges: (u, v) has the same colour as its endpoints, or is a blue edge in the special case where u is red and v is green. Our idea is to keep the colouring and flip all non-green edges, thus obtaining a single source graph H with source s. (Any vertex was either green and thus already reachable from s, or could reach some target t, and is reachable from s in H via the first green vertex on its path to t.)

Consider the single source reachability data structure for the red-green graph, H. This alone does not suffice to determine reachability in G, but it does when endowed with a few extra words per vertex:

M1 A red vertex u must remember the additional information of the best green vertices $BestGreen(u)$ on its own parent frame it can reach. There are at most 4 such vertices, one for each disegment.

M2 Information about paths from a red to a green vertex in the same component. See Section 4.1.

M3 Information about paths from a red vertex in some component C to a green vertex in an ancestor component of C. See Section 4.2.

Given a green vertex v, we know for each ancestral frame segment the best vertex that can reach v. For a red vertex u, given a segment p on an ancestral frame to u, we have information about the best vertex on p that may reach u in H via "ingoing" edges, that is, an edge from the corresponding $\widehat{F}_i(u)$. If that best vertex is red, then it is the best vertex on p that u can reach, again, from the "inside".

We may now case reachability based on the colour of nodes:

- For green u and red v, $\text{reach}_G(u, v)$ = No.
- For green vertices u, v, $\text{reach}_G(u, v) = \text{reach}_H(u, v)$
- For red vertices u, v, $\text{reach}_G(u, v) = \text{reach}_H(v, u)$
- When u is red and v is green, to determine $\text{reach}_G(u, v)$ we need more work. It will depend on where in the hierarchy of components, u and v reside.

When u is red and v is green, there are the following cases.

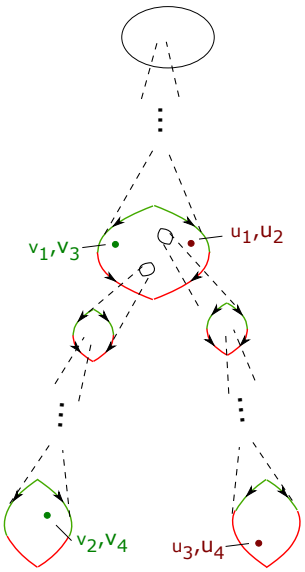

1. $c[u] = c[v]$. There may be a path from u to v:
   - Via a green vertex w in the parent frame of u. For each candidate $w \in BestGreen(u)$, try $\text{reach}_H(w, v)$. (See M1).
   - Staying within the frame, that is, $\text{reach}_{c[u]}(u, v)$. To handle this case we need to store more information, see Section 4.1.
2. $c[u] \prec c[v]$. There may be a path from u to v:
   - Via a green vertex w in the parent frame of u, $\text{reach}_H(w, v)$. (See M1).
   - Via a green vertex w, where $c[w] = c[u]$, then $\text{reach}_G(u, w)$ is in case 1 above. v knows the at most 4 such ws from the single source structure.
3. $c[u] \succ c[v]$. There may be a path from u to v:
   - Via a red edge (w', w) in G with $c[w] \preceq c[v] \prec c[w'] \preceq c[u]$. That is, in the single-source structure for H, u can find its best vertex w for each disegment of the parent frame of $c[v]$. For a path via that disegment to exist, w must be red, and $\text{reach}_G(w, v)$, which is in case 1 or 2 above, must return true.
   - Via a blue edge (w', w) with $c[w] \preceq c[v] \prec c[w'] \preceq c[u]$. We handle this case in Section 4.2.
4. $c[u], c[v] \succ N$, where $N = \text{lca}(c[u], c[v])$. A path from u to v must go:
   - Via w, $c[w] \preceq N$, then $\text{reach}_G(u, w)$ is in case 3 above. v computes at most 4 such ws from the single source structure, and note that all the vertices that v computes must be green.

4.1 Intracomponental blue edges

Consider the set of "blue" edges (a, b) from G where both the red vertex a and green b reside in some given component in the s-t-decomposition of H.

Lemma 4.1. We may assign to each vertex ≤ 2 numbers, such that if red u remembers $i, j \in \mathbb{N}$ and green v remembers $l, r \in \mathbb{N}$, then u can reach v if and only if $i \leq l \leq j$ or $i \leq r \leq j$ or $\min\{l, r\} \leq j < i$ or $j < i \leq \max\{l, r\}$.

Proof. The key observation is that we may enumerate all blue edges $b_0 = (u_0, v_0), \ldots b_i = (u_m, v_m)$ such that any red vertex can reach a segment of their endpoints, $v_i, \ldots, v_j$. Namely, the blue edges form a minimal cut in the planar graph which separates the red from the green vertices, and this cut induces a cyclic order. In this order, each red vertex may reach a segment of blue edges, and each green vertex may reach a segment of blue edge endpoints. Thus, the blue edge endpoints reachable from a given red vertex (through any path) is a union of overlapping segments, which is again a segment.

Now each red vertex remembers the indices of the first v_i and last v_j blue edge endpoint it may reach. For a green vertex v, the s-t-subgraph with v as target has a delimiting face consisting of two paths, P and Q. v remembers the indices l, r of the latest blue edge endpoints $v_l \in P$ and $v_r \in Q$, if they exist. Clearly, if l or r is within range, u may reach v. Contrarily, if u may reach v, it must do so via some vertex v' on $P \cup Q$. But then v' must be able to reach v_l or v_r, and thus, l or r is within range. □

4.2 Intercomponental blue edges

For any red vertex u, if a blue edge (u', v) reachable from u is separated u by a frame, then one of the best red vertices on that frame can reach u'. So let each red vertex remember the best ≤ 4 blue edges it can reach on its own frame. Then we can define 4 bitmasks $\{B^\beta(u)\}_{0 \leq \beta \leq 3}$ such that for any i finding the highest 1-bit $\leq i$ in each, gives at most 4 levels such that the best red vertices reachable from u on those levels together know the best blue edges for u.

References

[1] S. Alstrup, J. Holm, K. de Lichtenberg, and M. Thorup. Maintaining information in fully dynamic trees with top trees. *ACM Trans. Algorithms*, 1(2):243–264, October 2005.

[2] S. Alstrup, J. P. Secher, and M. Spork. Optimal on-line decremental connectivity in trees. *Inf. Process. Lett.*, 64:161–164, 1997.

[3] S. Arikati, D.Z. Chen, L.P. Chew, G. Das, M. Smid, and C.D. Zaroliagis. Planar spanners and approximate shortest path queries among obstacles in the plane. In *ESA '96*, pages 514–528, 1996.

[4] D.Z. Chen and J. Xu. Shortest path queries in planar graphs. In *STOC '00*, pages 469–478, 2000.

[5] H. Djidjev. Efficient algorithms for shortest path queries in planar digraphs. In *WG '96*, pages 151–165, 1996.

[6] H. Djidjev, G. Panziou, and C. Zaroliagis. Computing shortest paths and distances in planar graphs. In *ICALP '91*, pages 327–339, 1991.

[7] H. Djidjev, G. Panziou, and C. Zaroliagis. Fast algorithms for maintaining shortest paths in outerplanar and planar digraphs. In *FCT '95*, pages 191–200, 1995.

[8] D. Harel and R. Tarjan. Fast algorithms for finding nearest common ancestors. *SIAM J. Comput.*, 13(2):338–355, 1984.

[9] T. Kameda. On the vector representation of the reachability in planar directed graphs. *Inf. Process. Lett.*, 3(3):75–77, 1975.

[10] K. Kawarabayashi, P.N. Klein, and C. Sommer. Linear-space approximate distance oracles for planar, bounded-genus, and minor-free graphs. In *ALP '11*, pages 135–146, 2011.

[11] K. Kawarabayashi, C. Sommer, and M. Thorup. More compact oracles for approximate distances in undirected planar graphs. In *SODA '13*, pages 550–563, 2013.

[12] B.W. Kernighan and D.M. Ritchie. *The C Programming Language*. Prentice Hall, 2nd edition, 1988.

[13] P. Klein. Preprocessing an undirected planar network to enable fast approximate distance queries. In *SODA '02*, pages 820–827, 2002.

[14] P. B. Miltersen. Lower bounds for static dictionaries on rams with bit operations but no multiplication. In *ICALP '96*, pages 442–453. 1996.

[15] S. Mozes and C. Sommer. Exact distance oracles for planar graphs. In *SODA '12*, pages 209–222, 2012.

[16] M. Pătraşcu. Unifying the landscape of cell-probe lower bounds. *SIAM J. Comput.*, 40(3):827–847, 2011. Announced at FOCS'08. See also arXiv:1010.3783.

[17] R. Tamassia and I.G. Tollis. Dynamic reachability in planar digraphs with one source and one sink. *Theor. Comput. Sci.*, 119(2):331–343, 1993.

[18] R. Tarjan. Depth first search and linear graph algorithms. *SIAM J. Comput.*, 1972.

[19] M. Thorup. Compact oracles for reachability and approximate distances in planar digraphs. *J. ACM*, 51(6):993–1024, 2004.

[20] M. Thorup and U. Zwick. Approximate distance oracles. *J. ACM*, 52(1):183–192, 2005. Announced at STOC'01.